IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	§	Group Art Unit: 2152
	§	
Kuldipsingh A. Pabla Steven Waterhouse	§	Examiner: Hoang, Hieu T.
	§	
	§	Atty. Dkt. No.: 5681-08300 P7118
Serial No. 10/657,976	§	
	§	
	§	
Filed: September 9, 2003	§	
	§	
For: Peer-to-Peer Content Sharing/Distribution Networks	§	
	§	
	§	

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed December 20, 2007, Appellant presents this Appeal Brief. Appellant respectfully requests that the Board of Patent Appeals and Interferences consider this appeal.

I. REAL PARTY IN INTEREST

As evidenced by the assignment recorded at Reel/Frame 011070/0082, the subject application is owned by Sun Microsystems, Inc., a corporation organized and existing under and by virtue of the laws of the State of Delaware, and now having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054.

II. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences or judicial proceedings are known which would be related to, directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-35 are pending and stand finally rejected. The rejection of claims 1-35 is being appealed. A copy of claims 1-35 as currently pending is included in the Claims Appendix herein below.

IV. STATUS OF AMENDMENTS

No amendments have been submitted subsequent to the present final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a system comprising a plurality of peer nodes coupled to a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8; Figures 34A and 34B; page 16, lines 24-27. At least one of the plurality of peer nodes may be configured as a publisher peer node for one or more contents cached on the peer node. *See, e.g.,* page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24. A publisher peer node may be, for example, a primary, secondary, or an edge publisher peer node; *see, e.g.,* page 12, lines 22-26; page 15, lines 1-4. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12.

Each publisher peer node may be configured to publish one or more advertisements on the network. *See, e.g.,* page 15, lines 7-9; page 16, lines 27-29; page 79, lines 13-29; Figure 29A; page 81, line 25-page 82, line 16. Each advertisement may correspond to one of the one or more contents cached on the peer node. *See, e.g.,* Figure 10; page 100, line 1 – page 101, line 15. Each advertisement includes information for requesting a corresponding content. *See, e.g.,* page 15, lines 21-28; Figure 35; page 17, lines 13-20; page 83, lines 12-14; Figure 10; page 100, line 1 – page 101, line 15.

At least a subset of the plurality of peer nodes may each be configured to discover published advertisements on the network. *See, e.g.,* page 15, lines 14-26; Figure 35; page 17, lines 13-18; Figure 30; page 83, lines 1-14. The peers may also be configured to request content corresponding to the discovered advertisements in accordance with the information included in the advertisements. *See, e.g.,* page 15, lines 21-28; Figures 34A and 34B; page 16, line 28 - page 17, line 2; Figure 35; page 17, lines 13-20; Figure 37, element 504; page 18, lines 8-14.

A publisher peer node that caches a content corresponding to a discovered advertisement may be configured to provide the content corresponding to the discovered advertisement to a requesting peer node in response to a request for the content from the

requesting peer node. *See, e.g.,* page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14.

The requesting peer node may be configured to cache the content and become a content publisher peer node for the content corresponding to the discovered advertisement. *See, e.g.,* page 14, lines 9-12; page 15, line 28 – page 16, line 1; Figure 34B; page 17, lines 1-4; Figure 37, elements 506 and 507; page 18, lines 14-17.

Independent claim 8 is directed to a system comprising a plurality of content publisher peer nodes coupled to a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8; page 15, lines 1-12; Figure 34B; page 17, lines 1-4. Each of the plurality of content publisher peer nodes may be configured to cache user-requestable contents and to publish the cached contents on the network. *See, e.g.,* page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12. A content publisher peer node may be, for example, a primary, secondary, or an edge publisher peer node; *see, e.g.,* page 12, lines 22-26; page 15, lines 1-4.

The system may also comprise a content consumer peer node coupled to the network. *See, e.g.,* page 12, lines 6-8 and lines 22-23; Figures 34A and 34B, peers 400B and 400C; page 17, lines 1-2. The content consumer peer node may be configured to send a request for a particular content on the network in response to a user request for the particular content. *See, e.g.,* page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14; Figure 38, element 520; page 18, lines 26-27.

The content consumer peer node may receive the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network. *See, e.g.,* page 15, lines 14-19; Figure 35; page 17, lines 10-12; Figure 38, element 522; page 18, line 27 - page 19, line 10.

A logically nearest peer node is a peer node to which communications over the network take the least time. *See, e.g.,* page 13, lines 4-10.

Independent claim 12 is directed to a system, comprising a primary content publisher peer node configured to cache user-requestable contents and publish the cached contents for access by other peer nodes on a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8 and lines 22-26; page 15, lines 1-12; Figure 34B; page 17, lines 1-4; page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12.

The system may also comprise an edge content publisher peer node. *See, e.g.,* page 12, lines 22-26; page 15, lines 1-4; page 15, line 28 – page 16, line 2. The edge content publisher peer node may be configured to receive the user-requestable contents from the primary content publisher peer node. *See, e.g.,* page 16, lines 4-7 and lines 14-15. The edge content publisher peer node may be configured to cache the received contents and publish the received contents for access by the other peer nodes on the network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8 and lines 22-26; page 15, lines 1-12; Figure 34B; page 17, lines 1-4; page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12.

Independent claim 18 is directed to a system comprising means for a plurality of peer nodes to cache user-requestable contents and publish the user-requestable contents for access by other peer nodes on a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8; page 15, lines 1-12; page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24; page 17, lines 1-4. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12. A content publisher peer node may be, for example, a primary, secondary, or an edge publisher peer node; *see, e.g.,* page 12, lines 22-26; page 15, lines 1-4.

The system may also comprise means for a peer node to send a request for a particular content on the network in response to a user request for the particular content. *See, e.g.,* page 12, lines 6-8 and lines 22-23; Figures 34A and 34B, peers 400B and 400C; page 17, lines 1-2; page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14; Figure 38, element 520; page 18, lines 26-27.

The system may also comprise means for the peer node to receive the requested particular content from a nearest one of the plurality of peer nodes that caches and publishes the particular content on the network. *See, e.g.,* page 15, lines 14-19; Figure 35; page 17, lines 10-12; Figure 38, element 522; page 18, line 27 - page 19, line 10; page 13, lines 4-10.

Independent claim 20 is directed to a method comprising a content publisher peer node caching user-requestable contents and publishing the cached user-requestable contents for access by other peer nodes on a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8; page 15, lines 1-12; page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24; page 17, lines 1-4. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12. A content publisher peer node may be, for example, a primary, secondary, or an edge publisher peer node; *see, e.g.,* page 12, lines 22-26; page 15, lines 1-4.

The method also comprises one of the other peer nodes requesting a particular content on the network in response to a user request for the particular content. *See, e.g.,* page 12, lines 6-8 and lines 22-23; Figures 34A and 34B, peers 400B and 400C; page 17, lines 1-2; page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14; Figure 38, element 520; page 18, lines 26-27.

The method also comprises the one of the other peer nodes receiving the particular content from the content publisher peer node. *See, e.g.,* page 15, lines 14-18

and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14.

The method also comprises the one of the other peer nodes caching the received particular content and publishing the received particular content for access by the other peer nodes on the network. *See, e.g.,* page 14, lines 9-12; page 15, line 28 – page 16, line 1; Figure 34B; page 17, lines 1-4; Figure 37, elements 506 and 507; page 18, lines 14-17.

Independent claim 28 recites a storage medium including program instructions that are computer-executable to implement the method described above regarding claim 20. The program instructions may be computer-executable to implement a content publisher peer node caching user-requestable contents and publishing the cached user-requestable contents for access by other peer nodes on a network. *See, e.g.,* Figures 1A and 1B; page 2, lines 19-24; page 5, lines 17-22; page 12, lines 3-8; page 15, lines 1-12; page 11, lines 12-14; Figures 34A and 34B, peer 400A; page 16, lines 22-24; page 17, lines 1-4. For *content*, *see, e.g.* page 11, lines 16-20. For cache, *see, e.g.,* page 15, lines 9-12. A content publisher peer node may be, for example, a primary, secondary, or an edge publisher peer node; *see, e.g.,* page 12, lines 22-26; page 15, lines 1-4.

The program instructions may be computer-executable to implement one of the other peer nodes requesting a particular content on the network in response to a user request for the particular content. *See, e.g.,* page 12, lines 6-8 and lines 22-23; Figures 34A and 34B, peers 400B and 400C; page 17, lines 1-2; page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14; Figure 38, element 520; page 18, lines 26-27.

The program instructions may be computer-executable to implement the one of the other peer nodes receiving the particular content from the content publisher peer node. *See, e.g.,* page 15, lines 14-18 and lines 26-28; Figure 35; page 17, lines 6-12; Figure 37, element 504; page 18, lines 8-14.

The program instructions may be computer-executable to implement the one of the other peer nodes caching the received particular content and publishing the received particular content for access by the other peer nodes on the network. *See, e.g.,* page 14, lines 9-12; page 15, line 28 – page 16, line 1; Figure 34B; page 17, lines 1-4; Figure 37, elements 506 and 507; page 18, lines 14-17.

The summary above describes various examples and embodiments of the claimed subject matter; however, the claims are not necessarily limited to any of these examples and embodiments. The claims should be interpreted based on the wording of the respective claims.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1, 4, 12, 15, 20, 24, 25, 28, 32 and 33 stand finally rejected under 35 U.S.C. § 102(e) as being anticipated by Saulpaugh et al. (U.S. Publication 2004/0122903) (hereinafter "Saulpaugh").

2. Claims 2, 3, 5, 13, 14, 21-23 and 29-31 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of in view of Marmor et al. (U.S. Publication 2002/0062310) (hereinafter "Marmor") and Leber et al. (U.S. Publication 2003/0233455) (hereinafter "Leber").

3. Claims 8, 9, 18 and 19 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of Leber.

4. Claims 6, 7, 16, 17, 26, 27, 34 and 35 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of Lehikoinen et al. (U.S. Publication 2004/0260701) (hereinafter "Lehikoinen").

5. Claims 10 and 11 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh and Leber and further in view of Lehikoinen.

VII. ARGUMENT

First Ground of Rejection

Claims 1, 4, 12, 15, 20, 24, 25, 28, 32 and 33 stand finally rejected under 35 U.S.C. § 102(e) as being anticipated by Saulpaugh et al. (U.S. Publication 2004/0122903) (hereinafter "Saulpaugh"). Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claim 1

1. Saulpaugh's roles and messages are not analogous to content as recited in the context of claim 1 of the instant application.

Saulpaugh's "role" is clearly defined as a "location-independent address for a computer network" (paragraph [0068]). Saulpaugh's definition of "role" is clearly not analogous to "content" as recited in claim 1 of the instant application.

In the Action mailed September 26, 2007, the Examiner refers to FIG. 2 and paragraphs [0007] and [0056] of Saulpaugh. The Examiner appears to argue that, because Saulpaugh discloses in paragraph [0056] "in various embodiments any of various kinds of client application software 128 may utilize the T&R layer software 130 to send and receive messages for any desired purpose," and that "given the broadest reasonable interpretation, content can mean application data," that therefore "Saulpaugh discloses a role is a unique identifier used to identify application data or resources – which can be, e.g., a file that needs to be shared using a P2P file sharing application." Appellants respectfully traverse this argument.

Contrary to the Examiner's assertion, even under the Examiner's reasoning, Saulpaugh does not teach or suggest that "a role is a unique identifier used to identify application data or resources - which can be, e.g., a file that needs to be

shared using a P2P file sharing application." Saulpaugh clearly does not disclose or suggest that roles are or may be used to "identify application data or resources – which can be, e.g., a file that needs to be shared." Instead, Saulpaugh specifically discloses that "roles" are location-independent addresses used to identify sets of nodes on a network in Saulpaugh's messaging scheme, so that a node can send a message to a set of nodes using the "role" as an address instead of addressing each node individually or broadcasting to all nodes on the network. This is made clear in paragraph [0008] (emphasis added):

> The plurality of nodes may include a first node which is operable to send a message addressed using a "role". The role may be associated with one or more other nodes coupled to the network. The message may be sent to each of the one or more nodes with which the role is associated without specifying locations of the one or more nodes. For example, the role may have an associated name, and the message may be addressed using the role name (and possibly other information, such as a tree or subset of nodes with which the role is associated).

Note that substituting "application data or resources" or "a file" in the above for "role" or "node" would make no sense. "Application data or resources" or "files", or content, do not send messages to other "application data or resources" "files", or content on the network. Nodes send and receive messages; not content. Nodes, or roles, as used by Saulpaugh are not content and cannot be content under any reasonable interpretation of Saulpaugh.

Furthermore, simply because Saulpaugh suggests that "in various embodiments any of various kinds of client application software 128 may utilize the T&R layer software 130 to send and receive messages for any desired purpose" is clearly not enough to maintain the assertion that Saulpaugh discloses what is specifically recited in claim 1. Saulpaugh does not disclose or suggest the system as recited in claim 1 of the instant application.

Furthermore, the Examiner asserts "Saulpaugh discloses a role is a unique identifier used to identify application data or resources – which can be, e.g., a file that needs to be shared using a P2P file sharing application ([0068], client application software creates roles, [0069], a role address is associated with an application and a

protocol)." In [0068], Saulpaugh discloses "…message addressing is based on the concept of a "role". As used herein, a role may refer to a location-independent address for a computer network. The T&R layer may include an interface allowing client application software to create a role on one or more nodes on a tree." Saulpaugh is simply disclosing that client applications can create roles. **Again, Saulpaugh explicitly teaches and requires that roles are location-independent addresses for nodes, not for content.**

In [0069], Saulpaugh discloses (emphasis added):

> …a message address may also include information identifying a protocol ID. The protocol ID may be associated with a client application that utilizes the T&R layer. Multiple protocols may utilize the same tree. Thus, each message may be sent on a particular tree and, more particularly, to a particular set of nodes on the tree, i.e., the nodes having the specified role. As the message arrives to each node on the specified tree and having the specified role, the protocol ID may be used to determine which protocol on the node or which portion of client application **software** receives the message.

Nowhere in [0068]-[0069] does Saulpaugh teach or suggest the notion "a role is a unique identifier used to identify application data or resources – which can be, e.g., a file." In [0069], Saulpaugh discloses (again) that the "role" is used to send messages to particular sets of nodes, and further discloses that the protocol ID may be associated with a client application on a node and may be used to determine which protocol on the node or which portion of client application **software** receives the message. Again, Saulpaugh does not teach or suggest that messages are sent to content, application data or resources; Saulpaugh teaches messages are sent to nodes according to roles assigned to the nodes and possibly to client application software according to a protocol ID.

The Examiner goes on to assert "A peer node can publish a role (e.g., a file segment it cached for a file sharing application, [0084]-[0089], [0097], e.g. P2P file sharing protocol and application)." **Again, contrary to the Examiner's assertion, Saulpaugh does not teach or suggest a "role" is a "file segment [a node] cached for a file sharing application."** Indeed, the very notion that a "role" in Saulpaugh could be a

"file segment" is **contrary to Saulpaugh's actual teaching**. Saulpaugh teaches that a "role" is associated with a node, and a node is clearly not a "file segment", nor could it be nor would it make sense for a role to be a file segment in Saulpaugh's messaging system. File segments do not send messages to other file segments on a network, nor do file segments receive messages from other file segments on a network.

The Examiner goes on to assert "In response to role publishing or advertisement, a requesting node (one that receives advertisements) may send a request for a role ([0082]) and may receive responses from advertising nodes that have the roles." Paragraph [0082] states that a client application may be able to "request a role (sends a message to the current role, requesting to become that role)." Saulpaugh is referring to a client application requesting a role, and in [0082] it would make no sense to substitute something like a "file" or a "file segment" as the requestor (as the node or client application) or as the role. A client application requesting a role, as disclosed in [0082], does not in any way teach or suggest what the Examiner is asserting – that a "role" may be a "file segment" or some other content. A role is simply a location-independent addressing scheme that may be used by client applications to assume the role of a "node" and to send messages to other nodes and receive messages from other nodes using that role as a location-independent address.

The Examiner goes on to assert "the response may include data ([0063]) – reading on 'wherein each advertisement corresponds to one of the one or more contents cached on the peer node', 'request content corresponding to the discovered advertisements in accordance with the information included in the advertisements', and 'cache the content and become a content publisher peer node for the content corresponding to the discovered advertisement." **Contrary to the Examiner's assertion "the concept of a message response including data may be integrated in a sender to receiver back to sender protocol provided by the T&R layer" clearly does not "read on" the specific limitations as recited in claim 1.** Saulpaugh noting that message responses in the T&R layer may include "data" is clearly insufficient to teach the specific limitations in claim 1 that the Examiner asserts Saulpaugh "reads upon."

In the Advisory Action dated December 7, 2007, in section (1), the Examiner first asserts "Saulpaugh discloses that a client can use application software to send and receive messages for any purpose." This argument is essentially the same as the argument from section (6) of the Action of September 26, 2007, traversals of which are provided above. The Examiner cites some additional sections from Saulpaugh. Paragraph [0072] reads:

> Role-based addressing may also allow distributed software to run in a peer-to-peer manner. Nodes do not need to keep track of global state, such as knowing which other nodes are present on the network or which roles are bound to which nodes. A node may simply accomplish an operation by routing a message to a particular role, without needing to know which particular node or nodes have the role.

This citation clearly does not teach any of the limitations from claim 1, and adds nothing of substance to the Examiner's previous, erroneous arguments. **If anything, this citation, in teaching "[a] node may simply accomplish an operation by routing a message to a particular role, without needing to know which particular node or nodes have the role" actually teaches against what is recited in claim 1**, which is directed at advertising content for discovery on a network via advertisements, discovery of the advertisements, and requesting of content corresponding to discovered advertisements in accordance with the information included in the advertisements.

The Examiner also cites paragraphs [0063]-[0065]. Paragraph [0063] was previously cited, and a response to that paragraph is provided above. Paragraphs [0064]-[0065] simply provide more detail on sending messages in Saulpaugh's system, do not teach any of the limitations as recited in claim 1, and add nothing of substance to the Examiner's previous arguments.

In section (1) of the Advisory Action, the Examiner further asserts "Role publishing and role replying messages both contain a body portion that has data to be copied or transferred from nodes to nodes in a p2p network…" First, **the Examiner has improperly asserted that both Saulpaugh's "roles" and the "data to be copied or transferred" in the "body portion" of Saulpaugh's messages sent to roles are**

analogous to content as recited in claim 1. Saulpaugh's messages are described as being sent to roles ("[a] node may simply accomplish an operation by routing a message to a particular role, without needing to know which particular node or nodes have the role").

The differences between what Saulpaugh teaches in regards to these messages and what is recited in claim 1 is clear. Claim 1 is directed at advertising content for discovery on a network via published advertisements, discovery of the advertisements, and requesting of content corresponding to discovered advertisements in accordance with the information included in the advertisements. A publisher peer node that caches a content corresponding to a discovered advertisement provides the content corresponding to the discovered advertisement to a requesting peer node in response to a request for the content from the requesting peer node. Saulpaugh does not teach that the "data to be copied or transferred" in the "body portion" of Saulpaugh's messages is a request for content, nor does Saulpaugh teach that the "data" is content provided by a publisher peer node to a requesting peer node in response to a request for the content, nor does Saulpaugh teach that the request for content, or any message for that matter, was sent by the requesting peer node in accordance with the information included in a discovered advertisement for the content, nor does Saulpaugh teach the advertisement of "data to be copied or transferred," nor does Saulpaugh teach discovery of advertisements for "data to be copied or transferred."

The Examiner cites Saulpaugh, paragraphs [0289]-[0303], which is part of a section titled "Messaging." Paragraphs [0290]-[00303] describe a "send" function and inputs to the function. Paragraph [0290] reads: "Send--This function passes an array of bytes from sender to all nodes holding the specified role." Paragraphs [0291]-[0303] simply describe "inputs to the function". The citation actually further highlights the difference between what Saulpaugh teaches and what is recited in claim 1. Moreover, even if the "array of bytes" was considered analogous to content as recited in claim 1, Saulpaugh does not teach, for the "array of bytes" in a message sent using the "send" function, anything like what is recited in claim 1 in regards to advertising content,

discovering advertisements, and requesting content in accordance with the information included in the advertisements.

Paragraph [0297], "especially" cited by the Examiner, reads: "Body of Message (byte[ ]): A variable length array of bytes to be copied from sender to all receivers (role instances)." Again, this actually servers to further highlight the difference between what Saulpaugh teaches and what is recited in claim 1. *Even if* the "array of bytes" was considered analogous to content as recited in claim 1, Saulpaugh does not teach, for the "array of bytes" in a message sent using the "send" function, anything like what is recited in claim 1. Saulpaugh teaches that the "array of bytes" is to be copied from the sender to all receivers (role instances). Saulpaugh's messages are sent to roles without needing to know which particular node or nodes have the role ("[a] node may simply accomplish an operation by routing a message to a particular role, without needing to know which particular node or nodes have the role"). In contrast, claim 1 recites advertising content, discovering advertisements, a requesting node requesting content in accordance with the information included in a discovered advertisement, and a publisher peer node sending the specific content to the specific requesting node in response to the request for content.

The Examiner also "especially" cites paragraph [0303], which reads:

> Last Reply Flag (boolean): Indicates that this reply is the last one from this node. Any subsequent replies are not allowed. A last reply signals that per-node bookkeeping to route responses back to a sender is no longer needed (at the local node).

This paragraph clearly does not teach or suggest any of the limitations of claim 1.

In section (1) of the Advisory Action, the Examiner further asserts "It would be a reasonable interpretation to read the body of the send and reply messages as content to be distributed." Appellants strongly disagree. Again, even if the "body of the send and reply messages" was considered analogous to content as recited in claim 1, Saulpaugh does not teach, for the "body of the send and reply messages", anything like what is recited in claim 1.

2. Saulpaugh does not disclose a publisher peer node configured to publish one or more advertisements on the network, wherein each advertisement corresponds to one of the one or more contents cached on the peer node, and wherein each advertisement includes information for requesting a corresponding content.

The Examiner has erroneously equated Saulpaugh's "roles" with content as recited in claim 1 of the instant application as noted above under subheading (1), and even if "the body of the send and reply messages" is interpreted as "content to be distributed", Saulpaugh does not teach the limitations as recited in claim 1 regarding content for at least the reasons noted above under subheading (1). Thus, for at least the reasons under subheading (1), Saulpaugh does not teach this element as recited in claim 1. Specifically, Saulpaugh does not teach advertising content at all. Even if the "body of the send and reply messages" is considered as "content to be distributed", nowhere does Saulpaugh teach the advertisement of "content to be distributed."

3. Saulpaugh does not disclose at least a subset of the plurality of peer nodes each configured to discover published advertisements on the network and request content corresponding to the discovered advertisements in accordance with the information included in the advertisements.

The Examiner has erroneously equated Saulpaugh's "roles" with content as recited in claim 1 of the instant application as noted above under subheading (1), and even if "the body of the send and reply messages" is interpreted as "content to be distributed", Saulpaugh does not teach the limitations as recited in claim 1 regarding content for at least the reasons noted above under subheading (1). Thus, for at least the reasons under subheading (1), Saulpaugh does not teach this element as recited in claim 1. Specifically, Saulpaugh does not teach requesting content corresponding to the discovered advertisements in accordance with the information included in the advertisements. In contrast, Saulpaugh teaches that the "array of bytes" is to be copied

from the sender to all receivers (role instances). Saulpaugh's messages are sent to roles without needing to know which particular node or nodes have the role ("[a] node may simply accomplish an operation by routing a message to a particular role, without needing to know which particular node or nodes have the role").

4. Saulpaugh does not disclose wherein the requesting peer node is configured to cache the content and become a content publisher peer node for the content corresponding to the discovered advertisement.

The Examiner has erroneously equated Saulpaugh's "roles" with content as recited in claim 1 of the instant application as noted above under subheading (1), and even if "the body of the send and reply messages" is interpreted as "content to be distributed", Saulpaugh does not teach the limitations as recited in claim 1 regarding content for at least the reasons noted above under subheading (1).

5. Saulpaugh does not disclose the elements as recited in the claim when viewed as a whole.

Saulpaugh is directed at a system and method for location-independent message addressing for a computer network. In Saulpaugh's system, a plurality of nodes connected to a network may include a first node that is operable to send a message addressed using a "role". The role may be associated with one or more other nodes coupled to the network. The message may be sent to each of the one or more nodes with which the role is associated without specifying locations of the one or more nodes (i.e., by simply specifying the role). (Saulpaugh, Abstract). Saulpaugh's "role" is defined as a "location-independent address for a computer network" (paragraph [0068]). A role may be initially associated with a first node, and the role may later be associated with another node in addition to the first node. If a message is then addressed to the role, the message may be sent to the additional node as well as the first node using the name of the role. (paragraph [0009]). Saulpaugh discloses that roles may be "published" (paragraphs [0078]-[0089]), and that a client application (on a node) may be able to "request a role".

To request a role, a node sends a message to the current role requesting to "become that role". (Paragraph [0082]).

In contrast to Saulpaugh, which is directed at "location-independent message addressing for a computer network" using roles (location-independent addresses for a computer network), claim 1 of the instant application is directed at caching, advertising and distributing **content** on a network. Moreover, contrary to the Examiner's assertion, Saulpaugh's "roles" are not analogous to "content." Saulpaugh clearly does not disclose elements as recited in claim 1 when viewed as a whole.

6. The Examiner has improperly tried to combine distinctly different aspects of Saulpaugh in an attempt to support the assertion that Saulpaugh anticipates claim 1.

In section (1) of the Advisory Action, the Examiner further asserts "Therefore, role publishing and requesting read on caching content, advertising content, requesting for content, and becoming a content distributor." Appellants respectfully disagree. The Examiner has used various citations from different portions of Saulpaugh, citations which do not teach the limitations as recited in claim 1, and has improperly tried to combine distinctly different aspects of Saulpaugh from the citations, **specifically "roles" and the "body of the send and reply messages"**, in an attempt to support the assertion that Saulpaugh anticipates claim 1. As shown above, Saulpaugh clearly does not anticipate claim 1.

Claim 4

As this claim depends from independent claim 1, Appellants traverse this rejection for at least the reasons presented above regarding claim 1. Additional arguments in regard to this rejection are presented under the subheadings below.

1. Saulpaugh does not disclose wherein the at least a subset of the plurality of peer nodes are each configured to: send a request for a particular content on the network; receive a portion of the particular content from a first content publisher peer node that caches the particular content in response to the request; and receive another portion of the particular content from a second content publisher peer node that also caches the particular content in response to the request.

The Examiner cites Saulpaugh, [0187], FIG. 75, and paragraph [0076], and asserts "a content portion is an instance of a role that can be received from any nodes that host that instance…a portion of content can be retrieved at one node, other portions from other nodes." The Examiner's assertion makes no sense in light of what Saulpaugh actually teaches. While it is true that Saulpaugh discloses that a node may request and receive "roles", the notion that "a portion of content can be retrieved at one node, other portions from other nodes" makes absolutely no sense in the context of Saulpaugh's "roles". Saulpaugh's "roles" are clearly defined as "location-independent addresses for a computer network" (paragraph [0068]). Saulpaugh's definition of "role" is clearly not analogous to "content" as recited in claim 4 of the instant application.

Paragraph [0076] simply describes routing messages to role instances, and teaches nothing like what is recited in claim 4.

FIG. 75 illustrates a sender node sending a message to multiple receiver nodes ([0030]: "FIG. 75 illustrates an example in which a message is sent from a sender node to multiple receiver nodes"), not a receiver node receiving portions of particular content from multiple sender nodes. Paragraph [0187] further describes this process, and states, in regard to FIG. 76, that the sender node may receive response messages from the receiver nodes. The reply messages may be aggregated into one response message that the sender then receives, or alternatively the sender node may receive all the response messages. Messages, in this context, are not analogous to "content" as recited in claim 4. Moreover, messages, in this context, are not analogous to "roles" as used by Saulpaugh.

Moreover, the Examiner is improperly trying to combine distinctly different aspects of Saulpaugh from the citations, **specifically "roles" and "messages"**, in an attempt to support the assertion that Saulpaugh anticipates claim 4. Roles are part of Saulpaugh's "T&R layer" that supports messaging as is described in FIGs. 75 and 76 and in paragraph [0187]. Messages are sent to roles in Saulpaugh's system. For example, in FIG. 75, the "sender node 320" would send a message to a "role."

Claim 12

1. The Examiner is erroneously equating Saulpaugh's "roles" with content as recited in claim 12 of the instant application.

Saulpaugh's "role" is defined as a "location-independent address for a computer network" (paragraph [0068]). Saulpaugh's definition of "role" is clearly not analogous to "content" as recited in claim 12 of the instant application. The Examiner has erroneously equated Saulpaugh's "roles" with content. See subheading (1) under claim 1 for further discussion.

2. Saulpaugh does not disclose a primary content publisher peer node configured to cache user-requestable contents and publish the cached contents for access by other peer nodes on a network.

The Examiner cites Saulpaugh, paragraph [0086], [0078] and [0082], which describe the publishing of role instances, and the ability of clients to request a role. The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 12.

3. Saulpaugh does not disclose an edge content publisher peer node configured to: receive the user-requestable contents from the primary content publisher peer node.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 12.

In addition, the Examiner appears to equate Saulpaugh's use of "edge" in [0086] with edge in "edge content publisher peer node" as recited in the claim. Unlike claim 12, Saulpaugh's edges correspond to links, not nodes: "an edge may be created that maps upon the link over which the publish message was received" (paragraph [0086]). Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 12.

4. Saulpaugh does not disclose an edge content publisher peer node configured to: cache the received contents.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 12.

In addition, the Examiner cites paragraphs [0063] and [0069], and asserts "response with data to be utilized by an application identified by a protocol ID." These citations discuss "message responses", which may include data. Note that the "data" in these citations is distinctly different than the role instances the Examiner above cited as analogous to "content," and furthermore the messaging "supported by the T&R layer" described by Saulpaugh refers to messages sent to roles, and thus is a different aspect of Saulpaugh than the notion of role instances and role publishing (which are part of the "T&R" layer that supports the messaging). The Examiner is improperly citing distinctly

different aspects of Saulpaugh in an attempt to assert that Saulpaugh anticipates the limitations of claim 12.

5. Saulpaugh does not disclose an edge content publisher peer node configured to: publish the received contents for access by the other peer nodes on the network.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 12.

In addition, the Examiner cites [0074] and [0086], and asserts "role instances can be broadcasted to other nodes or edge nodes which will receive and publish the broadcast message to other further nodes." First, Saulpaugh does not teach "edge nodes." As noted above in (6), Saulpaugh's "edge" corresponds to a link, not a node.

Second, the Examiner is improperly citing distinctly different aspects of Saulpaugh in an attempt to assert that Saulpaugh anticipates the limitations of claim 12. In this case, the Examiner has switched back to the notion of "roles" for this limitation, whereas for the previous limitation, the Examiner relied upon "message responses", which may include data. Again, Saulpaugh's "roles" and "message responses" are clearly and distinctly different aspects of Saulpaugh's system. Roles are published, and nodes may request and assume roles. Messages, however, are sent to roles. This is made clear in paragraph [0008] (emphasis added):

> The plurality of nodes may include a first node which is operable to send a message addressed using a "role". The role may be associated with one or more other nodes coupled to the network. The message may be sent to each of the one or more nodes with which the role is associated without specifying locations of the one or more nodes.

6. Saulpaugh does not disclose the elements as recited in the claim when viewed as a whole.

Saulpaugh is directed at "location-independent message addressing for a computer network" using roles (location-independent addresses for a computer network). In contrast, claim 12 of the instant application is directed at caching, publishing, and accessing **user-requestable contents** on a network. See subheading (5) under claim 1 for further discussion.

Claim 15

As this claim depends from independent claim 12, Appellants traverse this rejection for at least the reasons presented above regarding claim 12. Additional arguments in regard to this rejection are presented under the subheadings below.

1. Saulpaugh does not disclose an edge peer node configured to: send a request for particular content on the network in response to a user request for the particular content; receive a portion of the particular content from the primary content publisher peer node in response to the request; receive a redirection to the edge content publisher peer node from the primary content publisher peer node; and receive another portion of the particular content from the edge content publisher peer node in response to the redirection.

The Examiner cites Saulpaugh, paragraph [0076]. Paragraph [0076] simply describes routing messages to role instances, teaches that a message sent to a "role", a "location-independent address," may be routed to multiple instances of that role, but teaches nothing like what is actually recited in claim 15. The Examiner asserts, in regard to [0076], "a peer that receives a query for instances of a role may host one or more instances and know redirecting routes to remaining instances; it responses to the query by returning the instances that it hosts together with routing information to other edges that host the remaining instances." The Examiner's interpretation of what is actually taught in the paragraph is not even close to what is actually taught. First, the message is not a "query for instances of a role"; the message is a message sent to a role. Second, a peer that receives the message does not respond by "returning the instances [of the role] that it

hosts together with routing information to other edges that host the remaining instances"; the peer forwards the message to the known instances. This continues "until each instance of the role receives the message."

Furthermore, as noted above in (3) for claim 12, Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 15.

Moreover, the Examiner appears to be again improperly trying to combine distinctly different aspects of Saulpaugh from the citations, **specifically "roles" and "messages"**, in an attempt to support the assertion that Saulpaugh anticipates claim 4. Roles are part of Saulpaugh's "T&R layer" that supports messaging. Messages are sent to roles in Saulpaugh's system.

Claims 20 and 28

1. The Examiner is erroneously equating Saulpaugh's "roles" with content as recited in claim 20 of the instant application.

Saulpaugh's "role" is defined as a "location-independent address for a computer network" (paragraph [0068]). Saulpaugh's definition of "role" is clearly not analogous to "content" as recited in claim 20 of the instant application. The Examiner has erroneously equated Saulpaugh's "roles" with content. See subheading (1) under claim 1 for further discussion.

2. Saulpaugh does not disclose a content publisher peer node caching user-requestable contents and publishing the cached user-requestable contents for access by other peer nodes on a network.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 20.

3. Saulpaugh does not disclose one of the other peer nodes requesting a particular content on the network in response to a user request for the particular content, receiving the particular content from the content publisher peer node, caching the received particular content, and publishing the received particular content for access by the other peer nodes on the network.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, Saulpaugh does not teach this element as recited in claim 20.

In addition, in the Action mailed September 26, 2007, the Examiner asserts that Saulpaugh "discloses requesting a role through a client application [0082]." Paragraph [0082] actually states that a client application may "request a role (sends a message to the current role, requesting to become that role)." The paragraph does not state "requesting a role through a client application"; it states that a client application may request to become a role. In other words, a client application may request to become or assume a "role" so that the application can send and receive messages according to the location-independent addressing scheme using roles as taught by Saulpaugh. As noted above, Saulpaugh clearly does not teach that a role is content, or that a role could be content. Indeed, as noted above, it would not make sense for a "role" to be "content" in Saulpaugh's system. Again, a role is a location-independent address for sending messages to nodes, and nodes are clearly not nor could they be content in Saulpaugh's messaging system. Paragraph [0082] makes that even more clear; a client application requests to become a role. Clearly, Saulpaugh is not referring to content as a "role".

4. Saulpaugh does not disclose the elements as recited in the claim when viewed as a whole.

Saulpaugh is directed at "location-independent message addressing for a computer network" using roles (location-independent addresses for a computer network).

In contrast, claim 20 of the instant application is directed at caching, publishing, and distributing **user-requestable contents** on a network. See subheading (5) under claim 1 for further discussion.

Claims 24 and 32

As these claims depend from independent claims 20 and 28, respectively, Appellants traverse these rejections for at least the reasons presented above regarding claims 20 and 28. Additional arguments in regard to these rejections are presented under the subheadings below.

1. Saulpaugh does not disclose a different peer node requesting the particular content on the network in response to a user request for the particular content; the different peer node receiving a portion of the particular content from the content publisher peer node in response to the request; the different peer node receiving a redirection to the one of the other peer nodes from the content publisher peer node; and the different peer node receiving another portion of the particular content from the one of the other peer nodes in response to the redirection.

In the rejection of claim 15, the Examiner cites Saulpaugh, paragraph [0076]. Paragraph [0076] simply describes routing messages to role instances, teaches that a message sent to a "role", a "location-independent address," may be routed to multiple instances of that role, but teaches nothing like what is actually recited in claim 15. The Examiner asserts, in regard to [0076], "a peer that receives a query for instances of a role may host one or more instances and know redirecting routes to remaining instances; it responses to the query by returning the instances that it hosts together with routing information to other edges that host the remaining instances." The Examiner's interpretation of what is actually taught in the paragraph is not even close to what is actually taught. First, the message is not a "query for instances of a role"; the message is a message sent to a role. Second, a peer that receives the message does not respond by "returning the instances [of the role] that it hosts together with routing information to

other edges that host the remaining instances"; the peer forwards the message to the known instances. This continues "until each instance of the role receives the message."

Moreover, the Examiner is again improperly trying to combine distinctly different aspects of Saulpaugh from the citations, **specifically "roles" and "messages"**, in an attempt to support the assertion that Saulpaugh anticipates claim 4. Roles are part of Saulpaugh's "T&R layer" that supports messaging. Messages are sent to roles in Saulpaugh's system.

Claims 25 and 33

As these claims depend from independent claims 24 and 32, respectively, Appellants traverse these rejections for at least the reasons presented above regarding claims 24 and 32. Additional arguments in regard to these rejections are presented under the subheadings below.

1. Saulpaugh fails to disclose wherein the content publisher peer node is a primary publisher of the particular content, and wherein the one of the other peer nodes is an edge publisher of the particular content.

The Examiner appears to equate Saulpaugh's use of "edge" in [0074] and [0086] with edge in "edge publisher" as recited in the claim. Unlike claim 24, Saulpaugh's edges correspond to links, not nodes: "an edge may be created that maps upon the link over which the publish message was received" (paragraph [0086]). Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 24. Thus, the Examiner's definition "an edge publisher is a peer that receives the instances advertised by the primary publisher and itself publishes the instances to other peers" is not supported by the citations.

Second Ground of Rejection

Claims 2, 3, 5, 13, 14, 21-23 and 29-31 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of in view of Marmor et al. (U.S. Publication 2002/0062310) (hereinafter "Marmor") and Leber et al. (U.S. Publication 2003/0233455) (hereinafter "Leber"). Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claims 2, 3, and 5

As these claims depend from independent claim 1, Appellants traverse these rejections for at least the reasons presented above under the first ground of rejection regarding claim 1. Additional arguments in regard to these rejections are presented under the subheadings below.

1. The Examiner asserts that Saulpaugh discloses an element that is not recited in claim 2.

In the Action dated September 26, 2007, the Examiner asserts "Saulpaugh discloses the plurality of peer nodes comprises an edge peer node ([0075], edge peer nodes pointing to other peer nodes via which the query message can be sent for a respective role instance)." Claim 2 does not recite an edge peer node.

2. The cited art does not disclose wherein the at least a subset of the plurality of peer nodes are each configured to: discover two or more advertisements published by two or more content publisher peer nodes to advertise a particular content cached on each of the two or more content publisher peer nodes.

The Examiner relies upon Marmor to teach this limitation, citing FIGs. 4 and 5 and paragraphs [0013] and [0046]. Marmor is "directed to providing ways to allow the

use of peer-to-peer computer networks for typical commercial transactions, i.e., transactions in which a merchant sells a product to a customer, as opposed to transactions in which individuals merely exchange computer files" (Abstract). Thus, the purpose of Marmor's system appears to be opposed to what is recited in claim 2, which could be said to be directed to the exchange of "computer files", as content cached on content publisher peer nodes in a network would include the notion of computer files. In Marmor's system, a user at a mobile or fixed device may query for geographical locations of particular physical merchandise (e.g., street addresses of stores or distributors). One or more databases may be searched for geographical locations in response to the query. One or more located geographical locations may be returned to the user device, which may then determine a physically nearby location of the physical merchandise ("product") relative to the current physical location of the user device. In contrast, claim 2 recites peer nodes configured to discover two or more advertisements published by two or more content publisher peer nodes to advertise a particular content cached on each of the two or more content publisher peer nodes. Clearly, Marmot does not teach the limitation as recited in the claim.

3. The cited art does not disclose wherein the at least a subset of the plurality of peer nodes are each configured to: determine one of the two or more content publisher peer nodes as logically nearest on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time; and request the particular content from the logically nearest content publisher peer node in accordance with the advertisement corresponding to the logically nearest content publisher peer node.

The Examiner relies upon Leber, citing Abstract, [0005], and [0045], last 5 lines. Appellants note that the Examiner's citations, except for the Abstract, appear to be incorrect; paragraph [0005] simply defines a "computer apparatus" and paragraph [0045] contains a single sentence that states "FIG. 1a is a schematic block diagram depicting the computer program modules of the File Retrieval Software Application." Examiner asserts "a peer node probes for actual QoS (delay, bandwidth, packet loss) to each peer

that advertises the role and selects peers with best QoS qualifications for setting up service." What Leber discloses is outlined in the Abstract, cited by the Examiner (emphasis added):

> The method involves sending a request for a file to the server computer; receiving back from the server an authentication code and a list of peer client computers that have the requested file or part of it; sending a request for the file to a subset of peer clients that yield the fastest download rate; receiving file data back from this subset of peer clients; reassembling the requested file using data sent by the peer clients; and checking the integrity and completeness of the reconstructed file by comparing a computed checksum of said reconstructed file with the authentication code.

Leber is directed at a "distributed file sharing system for fast transfer of data [received] from multiple computer data storage mediums connected by peer-to-peer connections through a computer network" (*see, e.g.*, [0030]).

Furthermore, Leber teaches, in paragraph [0033], that "the present invention eliminates the requirement for a user to download an entire file from a single source and instead provides a system and a method for the transfer of multiple parts of a file from a plurality of peer client computers, that can be reassembled into one file and checked for completeness and integrity by a predetermined authentication procedure." Thus, Leber appears to teach against what is recited in claim 2: a peer node configured to "request the particular content from the logically nearest content publisher peer node in accordance with the advertisement corresponding to the logically nearest content publisher peer node." In Leber's system, a requested particular "content" would not be received from a logically nearest peer node; Leber teaches that a file or parts of a file would be received from a plurality of peer client computers and reassembled on the receiving node.

In the Advisory Action dated December 7, 2007, in section (2), the Examiner asserts "Leber teaches requesting a content (which can be multiple pieces of content) from logically nearest nodes. Each piece of content of Leber reads on the claimed content, and would be received from a nearest peer." First, Leber does not teach "requesting a content from logically nearest nodes." Leber teaches that a request for a

file is sent to a server computer: "The method involves sending a request for a file to the server computer." (Abstract) Furthermore, as noted above, Leber teaches that parts of the requested file are requested and received from multiple peer client computers, a list of which is sent to the requestor by the server, and the parts are reassembled by the requestor: "receiving back from the server an authentication code and a list of peer client computers that have the requested file or part of it; sending a request for the file to a subset of peer clients that yield the fastest download rate; receiving file data back from this subset of peer clients; reassembling the requested file using data sent by the peer clients; and checking the integrity and completeness of the reconstructed file by comparing a computed checksum of said reconstructed file with the authentication code." (Abstract) In contrast, claim 2 recites a peer node determines one of two or more content publisher peer nodes [for which advertisements for the particular content were discovered] as logically nearest on the network and requests the particular content from the logically nearest content publisher peer node in accordance with the advertisement corresponding to the logically nearest content publisher peer node. Leber clearly does not teach what is recited in claim 2. In fact, Leber, in teaching "a distributed file sharing system and a method for providing fast download of data from multiple data storage mediums" (Abstract), actually teaches against what is recited in claim 2.

4. The Examiner has not provided a proper *prima facie* reason to combine the references.

The Examiner's reasoning to combine the references does not appear to be a reason that would be applicable to Saulpaugh's system for location-independent message addressing for a computer network. Saulpaugh is not directed at "providing peer-to-peer services", but is instead directed at location-independent message addressing for a computer network. Nothing in Saulpaugh suggests that there would be any advantage to finding a "peer with best QoS qualifications" from which to receive a "role", which is simply a location-independent address and not a file. It is not at all obvious as to how Leber's "distributed file sharing system for fast transfer of data from multiple computer data storage mediums connected by peer-to-peer connections through a computer

network" would be applicable in Saulpaugh's system. Saulpaugh's "roles" are clearly not received in pieces from multiple data storage mediums and reconstructed on the requesting node, nor would it make any sense to apply Leber's system to Saulpaugh. Furthermore, the Examiner's reasoning to combine the references does not appear to be a reason that would be applicable to Marmot's system "directed to providing ways to allow the use of peer-to-peer computer networks for typical commercial transactions, i.e., transactions in which a merchant sells a product to a customer, as opposed to transactions in which individuals merely exchange computer files" (Abstract). The Marmon reference appears to be directed at something that is opposed to what Leber and Saulpaugh appear to be directed at, as both Leber and Saulpaugh appear to be directed along the lines of "transactions in which individuals merely exchange computer files" and not at all directed to "providing ways to allow the use of peer-to-peer computer networks for typical commercial transactions, i.e., transactions in which a merchant sells a product to a customer."

In addition, Leber, in teaching a "method for providing fast download of data from multiple data storage mediums," appears to teach against what is recited in claim 2, as noted above in (3). Furthermore, combining the references, even if possible, would not produce anything like what is recited in claim 2. Such a combination, if possible, would appear to be something like Saulpaugh's system of using roles (location-independent addressing) to send messages to "roles", that incorporates Leber's "method for providing fast download of data from multiple data storage mediums," and Marmor's system "providing ways to allow the use of peer-to-peer computer networks for typical commercial transactions, i.e., transactions in which a merchant sells a product to a customer" (Abstract). Whatever such a combination might result in, the result would be nothing like what is recited in claim 2. Moreover, the Examiner's reason is merely conclusory.

Claims 13 and 14

As these claims depend from independent claims 12, Appellants traverse these rejections for at least the reasons presented above under the first ground of rejection regarding claim 12.

In addition, Appellants traverse these rejections for at least the reasons presented above under claim 2, subheadings (2), (3), and (4).

In addition, the Examiner appears to equate Saulpaugh's use of "edge", e.g. in [0075], with edge in "edge peer node" as recited in the claim. Unlike claim 13, Saulpaugh's edges correspond to links, not nodes: Saulpaugh's edges correspond to links, not nodes: "an edge may be created that maps upon the link over which the publish message was received" (paragraph [0086]). Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 13.

Claims 21 and 22

As these claims depend from independent claim 20, Appellants traverse these rejections for at least the reasons presented above under the first ground of rejection regarding claim 20. In addition, Appellants traverse these rejections for at least the reasons presented above under claim 2, subheadings (2), (3), and (4).

Claim 23

As this claim depends from claim 21, Appellants traverse this rejection for at least the reasons presented above regarding claim 21. In addition, the Examiner appears to equate Saulpaugh's use of "edge", e.g. in [0075], with edge in "edge peer node" as recited in the claim. Unlike claim 23, Saulpaugh's edges correspond to links, not nodes: "an edge may be created that maps upon the link over which the publish message was

received" (paragraph [0086]). Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 23.

Claims 29 and 30

As these claims depend from independent claim 28, Appellants traverse these rejections for at least the reasons presented above under the first ground of rejection regarding claim 28. In addition, Appellants traverse these rejections for at least the reasons presented above regarding claim 2.

Claim 31

As this claim depends from claim 29, Appellants traverse this rejection for at least the reasons presented above regarding claim 29. In addition, the Examiner appears to equate Saulpaugh's use of "edge", e.g. in [0075], with edge in "edge peer node" as recited in the claim. Unlike claim 29, Saulpaugh's edges correspond to links, not nodes: Saulpaugh's edges correspond to links, not nodes: "an edge may be created that maps upon the link over which the publish message was received" (paragraph [0086]). Saulpaugh's use of "edge" is clearly and distinctly different than the use of "edge" in claim 29.

Third Ground of Rejection

Claims 8, 9, 18 and 19 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of Leber. Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claims 8, 18 and 19

1. The Examiner is erroneously equating Saulpaugh's "roles" with content as recited in claim 8 of the instant application.

Saulpaugh's "role" is defined as a "location-independent address for a computer network" (paragraph [0068]). Saulpaugh's definition of "role" is clearly not analogous to "content" as recited in claim 8 of the instant application. The Examiner has erroneously equated Saulpaugh's "roles" with content. Furthermore, even if "the body of the send and reply messages" is interpreted as "content to be distributed", Saulpaugh does not teach the limitations as recited in claim 8. See subheading (1) under claim 1 for further discussion.

2. The cited art does not disclose a plurality of content publisher peer nodes coupled to a network, wherein each of the plurality of content publisher peer nodes is configured to cache user-requestable contents and to publish the cached contents on the network.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, the cited art does not teach this element as recited in claim 8.

3. The cited art does not disclose a content consumer peer node coupled to the network and configured to send a request for a particular content on the network in response to a user request for the particular content.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, the cited art does not teach this element as recited in claim 8.

4. The cited art does not disclose a content consumer peer node coupled to the network and configured to receive the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time.

The Examiner has erroneously equated Saulpaugh's "roles" with content, as noted above under subheading (1). Thus, for at least this reason, the cited art does not teach this element as recited in claim 8.

In addition, the Examiner asserts that "Leber discloses [receive the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time].", and cites Leber's abstract, [0005], and [0045] last 5 lines. Appellants note that the Examiner's citations, except for the Abstract, appear to be incorrect; paragraph [0005] simply defines a "computer apparatus" and paragraph [0045] contains a single sentence that states "FIG. 1a is a schematic block diagram depicting the computer program modules of the File Retrieval Software Application." Examiner asserts "a peer node probes for actual QoS (delay, bandwidth, packet loss) to each peer that advertises the role and selects peers with best QoS qualifications for setting up service." What Leber discloses is outlined in the Abstract, cited by the Examiner (emphasis added):

> The method involves sending a request for a file to the server computer; receiving back from the server an authentication code and a list of peer client computers that have the requested file or part of it; sending a request for the file to a subset of peer clients that yield the fastest download rate; receiving file data back from this subset of peer clients; reassembling the requested file using data sent by the peer clients; and checking the integrity and completeness of the reconstructed file by comparing a computed checksum of said reconstructed file with the authentication code.

Leber is directed at a "distributed file sharing system for fast transfer of data [received] from multiple computer data storage mediums connected by peer-to-peer connections through a computer network" (*see, e.g.*, [0030]).

Furthermore, Leber teaches, in paragraph [0033], that "the present invention eliminates the requirement for a user to download an entire file from a single source and instead provides a system and a method for the transfer of multiple parts of a file from a plurality of peer client computers, that can be reassembled into one file and checked for completeness and integrity by a predetermined authentication procedure." Thus, Leber appears to teach against what is recited in claim 8: "sending a request for a particular content on the network in response to a user request for the particular content; and receiving the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network." In Leber's system, a requested particular "content" would not be received from a logically nearest peer node; Leber teaches that a file or parts of a file would be received from a plurality of peer client computers and reassembled on the receiving node.

In the Advisory Action dated December 7, 2007, in section (2), the Examiner asserts "Leber teaches requesting a content (which can be multiple pieces of content) from logically nearest nodes. Each piece of content of Leber reads on the claimed content, and would be received from a nearest peer." First, Leber does not teach "requesting a content from logically nearest nodes." Leber teaches that a request for a file is sent to a server computer: "The method involves sending a request for a file to the server computer." (Abstract) Furthermore, as noted above, Leber teaches that parts of the requested file are requested and received from multiple peer client computers, a list of which is sent to the requestor by the server, and the parts are reassembled by the requestor: "receiving back from the server an authentication code and a list of peer client computers that have the requested file or part of it; sending a request for the file to a subset of peer clients that yield the fastest download rate; receiving file data back from this subset of peer clients; reassembling the requested file using data sent by the peer clients; and checking the integrity and completeness of the reconstructed file by

comparing a computed checksum of said reconstructed file with the authentication code." (Abstract) In contrast, claim 8 recites a content consumer peer node that sends a request for a particular content on the network in response to a user request for the particular content; and receives the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network. Leber clearly does not teach what is recited in claim 8. In fact, Leber, in teaching "a distributed file sharing system and a method for providing fast download of data from multiple data storage mediums" (Abstract), actually teaches against what is recited in claim 8.

5. The cited art does not disclose the elements as recited in the claim when viewed as a whole.

The Saulpaugh reference is directed at "location-independent message addressing for a computer network" using roles (location-independent addresses for a computer network). In contrast, claim 8 of the instant application is directed at caching, publishing, and distributing user-requestable **contents** on a network. As noted above under subheading (1) under claim 1, contrary to the Examiner's assertion, Saulpaugh does not teach or suggest that "roles" are or could be content or "application data." According to Saulpaugh, roles (location-independent addresses) may be requested by client applications. Paragraph [0082] actually states that a client application may "request a role (sends a message to the current role, requesting to become that role)." Thus, a client application can "become" or assume a role, and may receive messages from other "nodes" via Saulpaugh's location-independent addressing scheme that uses "roles". Saulpaugh does not teach that "content" or "application data" can "become a role." Saulpaugh does not teach or suggest what the Examiner asserts in regard to claim 8. See subheading (1) under claim 1 for further discussion.

Leber is directed at a "distributed file sharing system for fast transfer of data [received] from multiple computer data storage mediums connected by peer-to-peer connections through a computer network" (*see, e.g.*, [0030]). In Leber's system, a requested particular "content" would not be received from a logically nearest peer node;

Leber actually teaches that a file or parts of a file would be received from a plurality of peer client computers and reassembled on the receiving node. Leber thus appears to teach against what is recited in claim 8.

Thus, the cited art, alone or in combination, does not disclose the elements as recited in the claim when viewed as a whole.

6. The Examiner has not provided a proper *prima facie* reason to combine the references.

The Examiner's reasoning to combine Saulpaugh and Leber does not appear to be a reason that would be applicable to Saulpaugh's system for location-independent message addressing for a computer network. Saulpaugh is not directed at "providing peer-to-peer services", but is instead directed at location-independent message addressing for a computer network. Nothing in Saulpaugh suggests that there would be any advantage to finding a "peer with best QoS qualifications" from which to receive a "role", which is simply a location-independent address and not a file. It is not at all obvious as to how Leber's "distributed file sharing system for fast transfer of data from multiple computer data storage mediums connected by peer-to-peer connections through a computer network" would be applicable in Saulpaugh's system. Saulpaugh's "roles" are clearly not received in pieces from multiple data storage mediums and reconstructed on the requesting node, nor would it make any sense to apply Leber's system to Saulpaugh. In addition, Leber, in teaching a "method for providing fast download of data from multiple data storage mediums," appears to teach against what is recited in claim 8, as noted above in (5). Furthermore, combining the references, even if possible, would not produce anything like what is recited in claim 8. Such a combination, if possible, would appear to be something like Saulpaugh's system of using roles (location-independent addressing) to send messages to "roles", that incorporates Leber's "method for providing fast download of data from multiple data storage mediums," which is nothing like what is recited in claim 8. Moreover, the Examiner's reason is merely conclusory.

In the Advisory Action dated December 7, 2007, in section (2), the Examiner asserts "Furthermore, having established that Saulpaugh does disclose content distributing, the combination of Saulpaugh and Leber would have a high expectation of success (both in p2p network), and make the p2p system more efficient by requesting and distributing content among peer nodes." Again, contrary to the Examiner's assertion, the cited art does not teach all the limitations of claim 8. Furthermore, combining the references, even if possible, would not produce anything like what is recited in claim 8. Moreover, the cited art, alone or in combination, does not teach what is recited in claim 8 when viewed as a whole.

Claim 9

1. The Examiner has not presented a proper *prima facie* rejection of the claim.

In the Action dated September 26, 2007, the Examiner rejected claim 9 "for the same rationale as claim 1." Claim 1 is rejected under 35 U.S.C. § 102(e) as being anticipated by Saulpaugh. Claim 9 depends from claim 8, which is rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of Leber. Moreover, claim 9, in combination with claim 8, is of different scope than claim 1. Thus, the Examiner has not presented a proper rejection of claim 9.

2. As claim 9 depends from independent claim 8, Appellants traverse the rejection of claim 9 for at least the reasons presented above regarding claim 8.

Fourth Ground of Rejection

Claims 6, 7, 16, 17, 26, 27, 34 and 35 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh in view of Lehikoinen et al. (U.S. Publication 2004/0260701) (hereinafter "Lehikoinen"). Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their

respective subheadings.

Claim 6

As this claim depends from independent claim 1, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 1.

Claim 7

As this claim depends from claim 1, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 1. Additional arguments in regard to this rejection are presented under the subheadings below.

1. The cited art does not disclose a plurality of peer nodes configured to participate in a peer-to-peer environment on the network in accordance with one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

The Examiner relies upon Lehikoinen, cites paragraph [0038], and broadly asserts "a peer group for file sharing." However the Examiner improperly broadly cites a passage from the reference in asserting that the reference discloses the limitations of the claim and does not show where the reference teaches each and every limitation of the claim. For example, the Examiner has failed to show where the reference teaches a plurality of peer nodes configured to participate in a peer-to-peer environment on the network in accordance with one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other. As another example, the Examiner has failed to show where the reference teaches a plurality of peer nodes configured to participate in a peer-to-peer environment on the network in accordance with one or more

peer-to-peer platform protocols for enabling the plurality of peer nodes to cooperate with each other to form peer groups. That the Lehikoinen reference may disclose "a peer group for file sharing" is insufficient to establish that Lehikoinen discloses all the specific limitations of claim 7.

Claim 16

As this claim depends from independent claim 12, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 12.

Claim 17

As this claim depends from claim 12, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 12. In addition, Appellants traverse this rejection for at least the reasons given under claim 7, subheading (1).

Claim 26

As this claim depends from independent claim 20, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 20.

Claim 27

As this claim depends from claim 20, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 20. In addition, Appellants traverse this rejection for at least the reasons given under claim 7, subheading (1).

Claim 34

As this claim depends from independent claim 28, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 28.

Claim 35

As this claim depends from claim 28, Appellants traverse this rejection for at least the reasons presented above under the first ground of rejection regarding claim 28. In addition, Appellants traverse this rejection for at least the reasons given under claim 7, subheading (1).

Fifth Ground of Rejection

Claims 10 and 11 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Saulpaugh and Leber and further in view of Lehikoinen. Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claim 10

As this claim depends from independent claim 8, Appellants traverse this rejection for at least the reasons presented above under the third ground of rejection regarding claim 8.

Claim 11

As this claim depends from claim 8, Appellants traverse this rejection for at least the reasons presented above under the third ground of rejection regarding claim 8. In

addition, Appellants traverse this rejection for at least the reasons given under the fourth ground of rejection, claim 7, subheading (1).

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-35 was erroneous, and reversal of the Examiner's decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee of $510.00 and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-08300/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Appellant

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: March 20, 2008

VIII. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A system, comprising:

a plurality of peer nodes coupled to a network;

at least one of the plurality of peer nodes configured as a publisher peer node for one or more contents cached on the peer node, wherein each publisher peer node is configured to publish one or more advertisements on the network, wherein each advertisement corresponds to one of the one or more contents cached on the peer node, and wherein each advertisement includes information for requesting a corresponding content; and

at least a subset of the plurality of peer nodes each configured to:

discover published advertisements on the network; and

request content corresponding to the discovered advertisements in accordance with the information included in the advertisements;

wherein a publisher peer node that caches a content corresponding to a discovered advertisement is configured to provide the content corresponding to the discovered advertisement to a requesting peer node in response to a request for the content from the requesting peer node; and

wherein the requesting peer node is configured to cache the content and become a content publisher peer node for the content corresponding to the discovered advertisement.

2. The system as recited in claim 1, wherein the at least a subset of the plurality of peer nodes are each configured to:

discover two or more advertisements published by two or more content publisher peer nodes to advertise a particular content cached on each of the two or more content publisher peer nodes;

determine one of the two or more content publisher peer nodes as logically nearest on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time; and

request the particular content from the logically nearest content publisher peer node in accordance with the advertisement corresponding to the logically nearest content publisher peer node.

3. The system as recited in claim 2, wherein the at least a subset of the plurality of peer nodes are each further configured to cache the particular content and become a content publisher peer node for the particular content.

4. The system as recited in claim 1, wherein the at least a subset of the plurality of peer nodes are each configured to:

send a request for a particular content on the network;

receive a portion of the particular content from a first content publisher peer node that caches the particular content in response to the request; and

receive another portion of the particular content from a second content publisher peer node that also caches the particular content in response to the request.

5. The system as recited in claim 1, wherein the at least a subset of the plurality

of peer nodes are each configured to:

broadcast a request for a particular content on the network;

receive a response to the request from each of two or more content publisher peer nodes that cache the particular content;

determine a logically nearest one of the two or more content publisher peer nodes on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time; and

get the content from the logically nearest content publisher peer node.

6. The system as recited in claim 1, wherein the at least a subset of the plurality of peer nodes are member peers in a peer group.

7. The system as recited in claim 1, wherein the plurality of peer nodes is configured to participate in a peer-to-peer environment on the network in accordance with one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

8. A system, comprising:

a plurality of content publisher peer nodes coupled to a network, wherein each of the plurality of content publisher peer nodes is configured to cache user-requestable contents and to publish the cached contents on the network;

a content consumer peer node coupled to the network and configured to:

send a request for a particular content on the network in response to a user request for the particular content; and

receive the particular content from a logically nearest content publisher peer node of the plurality of content publisher peer nodes on the network, wherein a logically nearest peer node is a peer node to which communications over the network take the least time.

9. The system as recited in claim 8, wherein the content consumer peer node is configured to become a content publisher peer node for the particular content, wherein to become a content publisher peer node for the particular content, the content consumer peer node is configured to cache the particular content and publish the particular content for access by other content consumer peer nodes on the network.

10. The system as recited in claim 8, wherein the plurality of peer nodes are member peers in a peer group.

11. The system as recited in claim 8, wherein the plurality of peer nodes is configured to participate in a peer-to-peer environment on the network in accordance with one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

12. A system, comprising:

a primary content publisher peer node configured to cache user-requestable contents and publish the cached contents for access by other peer nodes on a network;

an edge content publisher peer node configured to:

receive the user-requestable contents from the primary content publisher peer node;

cache the received contents; and

publish the received contents for access by the other peer nodes on the network.

13. The system as recited in claim 12, further comprising an edge peer node configured to:

send a request for particular content on the network in response to a user request for the particular content;

if the edge content publisher peer node is logically nearer to the edge peer node on the network than the primary content publisher peer node, receive the particular content from the edge content publisher peer node; and

if the primary content publisher peer node is logically nearer to the edge peer node on the network than the edge content publisher peer node, receive the particular content from the primary content publisher peer node;

wherein a logically nearer peer node is the peer node to which communications over the network take the least time.

14. The system as recited in claim 13, wherein the edge peer node is further configured to become a content publisher peer node for the particular content, wherein to become a content publisher peer node for the particular content, the edge peer node is configured to cache the particular content and publish the particular content for access by the other peer nodes on the network.

15. The system as recited in claim 12, further comprising an edge peer node configured to:

send a request for particular content on the network in response to a user request for the particular content;

receive a portion of the particular content from the primary content publisher peer node in response to the request;

receive a redirection to the edge content publisher peer node from the primary content publisher peer node; and

receive another portion of the particular content from the edge content publisher peer node in response to the redirection.

16. The system as recited in claim 12, wherein the peer nodes are member peers in a peer group.

17. The system as recited in claim 12, wherein the peer nodes are configured to participate in a peer-to-peer environment on the network in accordance with one or more peer-to-peer platform protocols for enabling the peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

18. A system, comprising:

means for a plurality of peer nodes to cache user-requestable contents and publish the user-requestable contents for access by other peer nodes on a network;

means for a peer node to send a request for a particular content on the network in response to a user request for the particular content; and

means for the peer node to receive the requested particular content from a nearest one of the plurality of peer nodes that caches and publishes the particular content on the network.

19. The system as recited in claim 18, further comprising means for the peer node to cache and publish the particular content for access by other peer nodes on the network.

20. A method, comprising:

a content publisher peer node caching user-requestable contents and publishing the cached user-requestable contents for access by other peer nodes on a network;

one of the other peer nodes:

requesting a particular content on the network in response to a user request for the particular content;

receiving the particular content from the content publisher peer node;

caching the received particular content; and

publishing the received particular content for access by the other peer nodes on the network.

21. The method as recited in claim 20, further comprising:

a different peer node requesting the particular content on the network in response to a user request for the particular content;

if the one of the other peer nodes is logically nearer to the different peer node on the network than the content publisher peer node, the different peer node receiving the particular content from the one of the other peer nodes; and

if the content publisher peer node is logically nearer to the different peer node on the network than the one of the other peer nodes, the different peer node receiving the particular content from the content publisher peer node;

wherein a logically nearer peer node is the peer node to which communications over the network take the least time.

22. The method as recited in claim 21, further comprising the different peer node caching the particular content and publishing the particular content for access by the other peer nodes on the network.

23. The method as recited in claim 21, wherein the different peer node is an edge peer node.

24. The method as recited in claim 20, further comprising:

a different peer node requesting the particular content on the network in response to a user request for the particular content;

the different peer node receiving a portion of the particular content from the content publisher peer node in response to the request;

the different peer node receiving a redirection to the one of the other peer nodes from the content publisher peer node; and

the different peer node receiving another portion of the particular content from the one of the other peer nodes in response to the redirection.

25. The method as recited in claim 20, wherein the content publisher peer node is a primary publisher of the particular content, and wherein the one of the other peer nodes is an edge publisher of the particular content.

26. The method as recited in claim 20, wherein the peer nodes are member peers in a peer group.

27. The method as recited in claim 20, wherein the peer nodes are configured to participate in a peer-to-peer networking environment implemented in accordance with one or more peer-to-peer platform protocols for enabling peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

28. A computer-accessible storage medium, comprising program instructions, wherein the program instructions are computer-executable to implement:

a content publisher peer node caching user-requestable contents and publishing the cached user-requestable contents for access by other peer nodes on a network;

one of the other peer nodes:

requesting a particular content on the network in response to a user request for the particular content;

receiving the particular content from the content publisher peer node;

caching the received particular content; and

publishing the received particular content for access by the other peer nodes on the network.

29. The computer-accessible storage medium as recited in claim 28, wherein the program instructions are further computer-executable to implement:

a different peer node requesting the particular content on the network in response to a user request for the particular content;

if the one of the other peer nodes is logically nearer to the different peer node on the network than the content publisher peer node, the different peer node receiving the particular content from the one of the other peer nodes; and

if the content publisher peer node is logically nearer to the different peer node on the network than the one of the other peer nodes, the different peer node receiving the particular content from the content publisher peer node;

wherein a logically nearer peer node is the peer node to which communications over the network take the least time.

30. The computer-accessible storage medium as recited in claim 29, wherein the program instructions are further computer-executable to implement the different peer node caching the particular content and publishing the particular content for access by the other peer nodes on the network.

31. The computer-accessible storage medium as recited in claim 29, wherein the different peer node is an edge peer node.

32. The computer-accessible storage medium as recited in claim 28, wherein the program instructions are further computer-executable to implement:

a different peer node requesting the particular content on the network in response to a user request for the particular content;

the different peer node receiving a portion of the particular content from the content publisher peer node in response to the request;

the different peer node receiving a redirection to the one of the other peer nodes from the content publisher peer node; and

the different peer node receiving another portion of the particular content from the one of the other peer nodes in response to the redirection.

33. The computer-accessible storage medium as recited in claim 28, wherein the content publisher peer node is a primary publisher of the particular content, and wherein the one of the other peer nodes is an edge publisher of the particular content.

34. The computer-accessible storage medium as recited in claim 28, wherein the peer nodes are member peers in a peer group.

35. The computer-accessible storage medium as recited in claim 28, wherein the peer nodes are configured to participate in a peer-to-peer networking environment implemented in accordance with one or more peer-to-peer platform protocols for enabling peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share network resources in the peer-to-peer environment.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.